arXiv:1702.04031v1 [stat.ME] 14 Feb 2017

MAXIMUM LIKELIHOOD ESTIMATION IN GAUSSIAN MODELS UNDER TOTAL POSITIVITY

STEFFEN LAURITZEN, CAROLINE UHLER, AND PIOTR ZWIERNIK

ABSTRACT. We analyze the problem of maximum likelihood estimation for Gaussian distributions that are multivariate totally positive of order two (MTP_2). By exploiting connections to phylogenetics and single-linkage clustering, we give a simple proof that the maximum likelihood estimator (MLE) for such distributions exists based on $n \geq 2$ observations, irrespective of the underlying dimension. Slawski and Hein [33], who first proved this result, also provided empirical evidence showing that the MTP_2 constraint serves as an implicit regularizer and leads to sparsity in the estimated inverse covariance matrix, determining what we name the ML graph. We show that the maximum weight spanning forest (MWSF) of the empirical correlation matrix is a spanning forest of the ML graph. In addition, we show that we can find an upper bound for the ML graph by adding edges to the MSWF corresponding to correlations in excess of those explained by the forest. This also gives new theoretical results in the study of inverse M-matrices. We provide globally convergent coordinate descent algorithms for calculating the MLE under the MTP_2 constraint which are structurally similar to iterative proportional scaling. We conclude the paper with a discussion of signed MTP_2 distributions.

1. INTRODUCTION

Total positivity is a special form of positive dependence between random variables that became an important concept in modern statistics; see, e.g., [3, 8, 20]. This property (also called the MTP_2 property) appeared in the study of stochastic orderings, asymptotic statistics, and in statistical physics [14, 28]. Families of distributions with this property lead to many computational advantages [2, 11, 30]. In a recent paper [12], the MTP_2 property was studied in the context of graphical models and conditional independence in general. It was shown that MTP_2 distributions have desirable Markov properties. Our paper can be seen as a continuation of this work with a focus on Gaussian distributions.

A p-variate real-valued distribution with density f w.r.t. a product measure μ is *multivariate totally positive of order 2* (MTP_2) if the density satisfies

$$f(x)f(y) \leq f(x \wedge y)f(x \vee y).$$

A multivariate Gaussian distribution with mean μ and covariance matrix Σ is MTP_2 if and only if the concentration matrix $K := \Sigma^{-1}$ is an *M-matrix*, that is, $K_{ij} \leq 0$

Date: February 15, 2017.

2010 *Mathematics Subject Classification.* 60E15, 62H99, 15B48.

Key words and phrases. MTP_2 distributions, Attractive Gaussian Markov random field (GMRF); non-frustrated GRMF; Gaussian graphical model; inverse M-matrix; ultrametric.

for all $i \neq j$ or, equivalently, if all partial correlations are nonnegative. Such distributions were considered by Bølviken [5] and Karlin and Rinott [22]. Moreover, Gaussian graphical models, or Gaussian Markov random fields, were studied in the context of totally positive distributions in [26]. MTP_2 Gaussian graphical models were shown to form a sub-class of *non-frustrated* Gaussian graphical models, which themselves are a sub-class of *walk-summable* Gaussian graphical models. Efficient structure estimation algorithms for MTP_2 Gaussian graphical models were given in [1] based on thresholding covariances after conditioning on subsets of variables of limited size. Efficient learning procedures based on convex optimization were suggested by Slawski and Hein [33] and this paper is closely related to their approach; see also [4].

Throughout this paper, we assume that we are given n i.i.d. samples from $\mathcal{N}(\mu, \Sigma)$, where Σ is a positive definite matrix of size $p \times p$. Without loss of generality, we assume that $\mu = 0$ and we focus on the estimation of Σ. We denote the sample covariance matrix based on the n samples by S. Then the log-likelihood function is, up to additive and multiplicative constants, given by

$$\ell(K; S) \quad = \quad \log\det K - \mathrm{tr}(SK).$$

We denote the cone of real symmetric matrices of size $p \times p$ by $\mathbb{S}^p$, its positive definite elements by $\mathbb{S}^p_{\succ 0}$, and its positive semidefinite elements by $\mathbb{S}^p_{\succeq 0}$. Note that $\ell(K; S)$ is a strictly concave function of $K \in \mathbb{S}^p_{\succeq 0}$. Since M-matrices form a convex subset of $\mathbb{S}^p_{\succ 0}$, the optimization problem for computing the *maximum likelihood estimator* (MLE) for MTP_2 Gaussian models is a convex optimization problem. Slawski and Hein [33] showed that the MLE exists, i.e., the global maximum of this optimization problem is attained, when $n \geq 2$. This yields a drastic reduction from $n \geq p$ without the MTP_2 constraint. In addition, they provided empirical evidence showing that the MTP_2 constraint serves as an implicit regularizer and leads to sparsity in the concentration matrix K.

In this paper, we analyze the sparsity pattern of the MLE $\hat{K}$ under the MTP_2 constraint. For a $p \times p$ matrix K we let $G(K)$ denote the undirected graph on p nodes with an edge uv if and only if $K_{uv} \neq 0$. Our main result is Theorem 4.2 that shows that with probability one, the smallest *maximum weight spanning forest* (MWSF) of the correlation matrix is a subgraph of the maximum likelihood (ML) graph $G(\hat{K})$. This provides a theoretical justification for initializing stepwise model selection methods in the MWSF of the sample correlation matrix. In addition, in Proposition 4.4 we obtain a simple upper bound for the ML graph $G(\hat{K})$ by adding edges to the MWSF corresponding to empirical correlations in excess of those provided by the MWSF. We illustrate these results in the following example.

Example 1.1. We consider the `carcass` data that are discussed in [17] and can be found in the `R`-library `gRbase`. This data set contains measurements of the thickness of meat and fat layers at different locations on the back of a slaughter pig together with the lean meat percentage on each of 344 carcasses. For our analysis we do not consider the lean meat percentage, since, by definition, this variable should be negatively correlated with the fat variables and positively correlated with the meat variables and hence the joint distribution is unlikely to be MTP_2.

The sample correlation matrix for this data set is given by

$$S = \begin{array}{c} \begin{matrix} \text{Fat11} & \text{Meat11} & \text{Fat12} & \text{Meat12} & \text{Fat13} & \text{Meat13} \end{matrix} \\ \begin{pmatrix} 1.00 & 0.04 & 0.84 & 0.08 & 0.82 & -0.03 \\ 0.04 & 1.00 & 0.04 & 0.87 & 0.13 & 0.86 \\ 0.84 & 0.04 & 1.00 & 0.01 & 0.83 & -0.03 \\ 0.08 & 0.87 & 0.01 & 1.00 & 0.11 & 0.90 \\ 0.82 & 0.13 & 0.83 & 0.11 & 1.00 & 0.02 \\ -0.03 & 0.86 & -0.03 & 0.90 & 0.02 & 1.00 \end{pmatrix} \end{array} \begin{matrix} \text{Fat11} \\ \text{Meat11} \\ \text{Fat12} \\ \text{Meat12} \\ \text{Fat13} \\ \text{Meat13} \end{matrix}$$

and its inverse is

$$S^{-1} = \begin{array}{c} \begin{matrix} \text{Fat11} & \text{Meat11} & \text{Fat12} & \text{Meat12} & \text{Fat13} & \text{Meat13} \end{matrix} \\ \begin{pmatrix} 4.53 & 0.78 & -2.34 & -1.74 & -1.72 & 1.01 \\ 0.78 & 5.10 & -0.25 & -2.51 & -0.76 & -2.12 \\ -2.34 & -0.25 & 4.52 & 1.40 & -1.93 & -0.94 \\ -1.74 & -2.51 & 1.40 & 7.00 & -0.09 & -4.11 \\ -1.72 & -0.76 & -1.93 & -0.09 & 4.11 & 0.53 \\ 1.01 & -2.12 & -0.94 & -4.11 & 0.53 & 6.51 \end{pmatrix} \end{array} \begin{matrix} \text{Fat11} \\ \text{Meat11} \\ \text{Fat12} \\ \text{Meat12} \\ \text{Fat13} \\ \text{Meat13} \end{matrix}$$

This empirical covariance matrix is not MTP_2; the positive entries in S^{-1} are highlighted in red. The MLE under MTP_2 can be computed for example using `cvx` [16] in `matlab` or using one of the simple coordinate descent algorithms discussed in Section 2. The MLE of the correlation matrix, rounded to 2 decimals, is

$$\hat{\Sigma} = \begin{array}{c} \begin{matrix} \text{Fat11} & \text{Meat11} & \text{Fat12} & \text{Meat12} & \text{Fat13} & \text{Meat13} \end{matrix} \\ \begin{pmatrix} 1.00 & 0.10 & 0.84 & 0.09 & 0.82 & 0.09 \\ 0.10 & 1.00 & 0.11 & 0.87 & 0.13 & 0.86 \\ 0.84 & 0.11 & 1.00 & 0.09 & 0.83 & 0.09 \\ 0.09 & 0.87 & 0.09 & 1.00 & 0.11 & 0.90 \\ 0.82 & 0.13 & 0.83 & 0.11 & 1.00 & 0.11 \\ 0.09 & 0.86 & 0.09 & 0.90 & 0.11 & 1.00 \end{pmatrix} \end{array} \begin{matrix} \text{Fat11} \\ \text{Meat11} \\ \text{Fat12} \\ \text{Meat12} \\ \text{Fat13} \\ \text{Meat13} \end{matrix}$$

The entries of $\hat{\Sigma}$ that changed compared to the sample correlation matrix S are highlighted in blue[1]. The sparsity pattern of $\hat{K} = \hat{\Sigma}^{-1}$ is captured by the ML graph $G(\hat{K})$ shown in Figure 1(a).

Note that all edges corresponding to blue entries in $\hat{\Sigma}$ are missing in this graph. As we show in Proposition 2.2, this is a consequence of the KKT conditions. Consider now the maximum weight spanning forest of the complete graph with weights given by the entries of S. In this example, the spanning forest is a chain represented by the thick red edges in Figure 1(a). By Theorem 4.2 these edges form a spanning tree of the ML graph $G(\hat{K})$. by Proposition 4.4 the excess correlation graph (EC graph) yields an upper bound for the ML graph $G(\hat{K})$. In this example the EC graph identifies the location of all non-edges apart from the edge Fat11 $\sim$ Meat12.

[1]We note that $\hat{\Sigma}_{45} > S_{45}$; the entries appear equal only because of the 2-digit rounding.

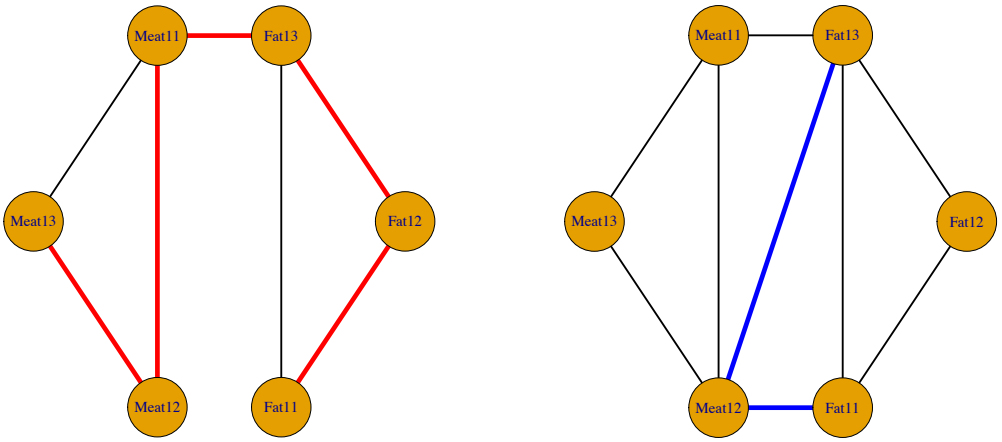

(a) MTP_2 constraint (b) Graphical lasso

FIGURE 1. Undirected Gaussian graphical models for `carcass` data. The thick red edges in (a) correspond to the MWSF of the correlation matrix and the blue edges in (b) to edges in the model chosen by the graphical lasso which are not in the ML graph of the MTP_2 solution.

Interestingly, applying various methods for model selection such as stepwise AIC, BIC, or graphical lasso all yield similar graphs. In Figure 1(b) we show the graphical model resulting from the graphical lasso, where the penalty parameter was chosen based on a robustness analysis. Note that the graph obtained under the MTP_2 constraint as well as the upper bound are subgraphs of the model obtained using the graphical lasso, indicating that ML estimation under the MTP_2 constraint might be an interesting alternative for model selection without the need of a tuning parameter. The extra edges in Figure 1(b) are represented by thick blue lines. □

The remainder of this paper is organized as follows: In Section 2, we review the duality theory that is known more generally for regular exponential families and specialize it to MTP_2 Gaussian distributions. This embeds the results by Slawski and Hein [33] into the framework of exponential families and also leads to two related coordinate descent algorithms for computing the MLE, one that acts on the entries of K and one that acts on the entries of Σ. In Section 3, we show how the problem of ML estimation for MTP_2 Gaussian distributions is connected to single-linkage clustering and ultrametrics as studied in phylogenetics. These observations result in a simple proof of the existence of the MLE for $n \geq 2$, a result that was first proven in [33]. Our proof is by constructing a primal and dual feasible point of the convex ML estimation problem for MTP_2 Gaussian models. In Section 4, we prove our main results, namely the characterization of a spanning forest of the ML graph $G(\hat{K})$ and a simple upper bound for it. Finally, in Section 5 we discuss how our results can be generalized to so-called *signed* MTP_2 Gaussian distributions, where the distribution is MTP_2 up to sign, i.e., $|X|$ is MTP_2. Such distributions were introduced by Karlin and Rinott in [21].

2. Duality theory for maximum likelihood estimation under MTP_2

We start this section by formally introducing MTP_2 distributions and then discuss the duality theory for Gaussian MTP_2 distributions. Let $V := \{1, 2, \ldots, p\}$ be a finite set and let $X = (X_v, v \in V)$ be a random vector with density f on the product space $\mathcal{X} = \prod_{v \in V} \mathcal{X}_v$, where $\mathcal{X}_v \subseteq \mathbb{R}$ is the state space of X_v, inheriting the order from $\mathbb{R}$. We define the coordinate-wise minimum and maximum as

$$x \wedge y = (\min(x_v, y_v), v \in V), \quad x \vee y = (\max(x_v, y_v), v \in V).$$

Then we say that X or the distribution of X is *multivariate totally positive of order two* (MTP_2) if its density function f on $\mathcal{X}$ satisfies

$$f(x)f(y) \quad \leq \quad f(x \wedge y)f(x \vee y) \qquad \text{for all } x, y \in \mathcal{X}. \tag{1}$$

In this paper, we concentrate on the Gaussian setting. It is easy to show that a Gaussian distribution with mean μ and covariance matrix Σ is MTP_2 if and only if $K = \Sigma^{-1}$ is an *M-matrix*, i.e.,

(i) $K_{vv} > 0$ for all $v \in V$,
(ii) $K_{uv} \leq 0$ for all $u, v \in V$ with $u \neq v$.

Properties of M-matrices were studied by Ostrowski [29] who chose the name to honor H. Minkowski. The connection to multivariate Gaussian distributions was established by Bølviken [5] and by Karlin and Rinott [22].

We denote the set of all M-matrices of size $p \times p$ by $\mathcal{M}^p$. Note that $\mathcal{M}^p$ is a convex cone. In fact, it is obtained by intersecting the positive definite cone $\mathbb{S}^p_{\succ 0}$ with all the coordinate half-spaces

$$\mathcal{H}^p_{ij} = \{X \in \mathbb{S}^p \mid X_{ij} \leq 0\}$$

with $i \neq j$. For a convex cone $\mathcal{C}$ we denote its closure by $\overline{\mathcal{C}}$. Then $\overline{\mathcal{M}^p}$ is given by $\mathbb{S}^p_{\succeq 0} \cap_{i<j} \mathcal{H}^p_{ij}$ and the ML estimation problem for Gaussian MTP_2 models can be formulated as the following optimization problem:

$$\begin{aligned} \underset{K}{\text{maximize}} \quad & \log\det(K) - \operatorname{trace}(KS) \\ \text{subject to} \quad & K \in \overline{\mathcal{M}^p} \end{aligned} \tag{2}$$

This is a convex optimization problem, since the objective function is concave on $\mathbb{S}^p_{\succeq 0}$.

Next, we introduce a second convex cone $\mathcal{N}^p$ that plays an important role for ML estimation in Gaussian MTP_2 models. To formally define this cone, we introduce two partial orders on matrices. Let A, B be two $p \times p$ matrices. Then $A \geq B$ means that $A_{ij} \geq B_{ij}$ for all $(i, j) \in V \times V$, and $A \succeq B$ means that $A - B \in \mathbb{S}^p_{\succeq 0}$. Then the cone $\mathcal{N}^p$ is defined as the negative closure of $\mathbb{S}^p_{\succ 0}$, i.e.

$$\mathcal{N}^p = \{X \in \mathbb{S}^p \mid \exists Y \in \mathbb{S}^p_{\succ 0} \text{ with } X \leq Y \text{ and } \operatorname{diag}(X) = \operatorname{diag}(Y)\}.$$

To simplify notation, we will suppress the dependence on p and write $\mathbb{S}$, $\mathbb{S}_{\succeq 0}$, $\mathbb{S}_{\succ 0}$, $\mathcal{M}$ and $\mathcal{N}$, when the dimension is clear. In the following result, we show that the cones $\mathcal{N}$ and $\mathcal{M}$ are dual to each other.

Lemma 2.1. *The closure of $\mathcal{N}$ is the dual to the cone of M-matrices $\mathcal{M}$, i.e.*

$$\overline{\mathcal{N}} = \big\{ S \in \mathbb{S} \mid \langle S, K\rangle \geq 0 \text{ for all } K \in \mathcal{M}\big\}. \tag{3}$$

Proof. We denote the dual of a convex cone $\mathcal{C}$ by $\mathcal{C}^\vee$. Let $\mathcal{C}_1, \mathcal{C}_2$ be two convex cones. Then it is an easy exercise to verify that

$$(\mathcal{C}_1 \cap \mathcal{C}_2)^\vee = \mathcal{C}_1^\vee + \mathcal{C}_2^\vee; \tag{4}$$

here $+$ denotes the Minkowski sum. Note that

$$\mathbb{S}_{\succ 0}^\vee = \mathbb{S}_{\succeq 0} \quad \text{and} \quad \mathcal{H}_{ij}^\vee = \mathcal{H}_{ij}.$$

This completes the proof, since $\mathcal{M} = \mathbb{S}_{\succ 0} \cap_{i<j} \mathcal{H}_{ij}$ and (4) can be applied inductively to any finite collection of convex cones. □

Using the cones $\mathcal{M}$ and $\mathcal{N}$ we now determine conditions for existence of the MLE in Gaussian MTP_2 models and give a characterization of the MLE. We say that the MLE does not exist if the likelihood does not attain the global maximum.

Proposition 2.2. *Consider a Gaussian* MTP_2 *model. Then the MLE $\hat{\Sigma}$ (and $\hat{K}$) exists for a given sample covariance matrix S on V if and only if $S \in \mathcal{N}$. It is then equal to the unique element $\hat{\Sigma} \succ 0$ that satisfies the following system of equations and inequalities*

$$(\hat{\Sigma}^{-1})_{uv} \leq 0 \text{ for all } u \neq v, \tag{5}$$

$$\hat{\Sigma}_{vv} - S_{vv} = 0 \text{ for all } v \in V, \tag{6}$$

$$(\hat{\Sigma}_{uv} - S_{uv}) \geq 0 \text{ for all } u \neq v, \tag{7}$$

$$(\hat{\Sigma}_{uv} - S_{uv})(\hat{\Sigma}^{-1})_{uv} = 0 \text{ for all } u \neq v, \tag{8}$$

Proof. The Lagrangian to (2) is given by

$$\mathcal{L}(K, \lambda) = \log\det(K) - \operatorname{trace}(KS) - \sum_{u<v} \lambda_{uv} K_{uv},$$

where we omitted the constraint $K \in \mathbb{S}_{\succ 0}$. This can be done, since it is assumed that K is in the domain of $\mathcal{L}$. Maximizing $\mathcal{L}$ with respect to K we obtain

$$(\hat{K}^{-1})_{uv} = \begin{cases} S_{vv} & \text{for } u = v \in V \\ S_{uv} + \lambda_{uv} & \text{otherwise.} \end{cases}$$

The Lagrange dual function is obtained by plugging in $\hat{K}$ for K in $\mathcal{L}$, i.e.,

$$g(\lambda) = \log\det(\hat{K}) - \operatorname{tr}(\hat{K}^{-1}\hat{K}) = \log\det(\hat{K}) - p.$$

Denoting $\hat{K}^{-1}$ by $\hat{\Sigma}$ and since $\lambda_{ij} \geq 0$, the dual optimization problem to (2) is given by

$$\begin{aligned} \underset{\Sigma \succeq 0}{\text{minimize}} \quad & -\log\det(\Sigma) - p \\ \text{subject to} \quad & \Sigma_{vv} = S_{vv}, \quad \text{for all } v \in V, \\ & \Sigma_{uv} \geq S_{uv}, \quad \text{for all } u \neq v, \end{aligned} \tag{9}$$

Note that the identity matrix is a strictly feasible point for (2). Hence, the MLE does not exist if and only if the likelihood is unbounded. Since by Slater's constraint

qualification strong duality holds for the optimization problems (2) and (9), the MLE does not exist if and only if $S \notin \mathcal{N}$.

Next we characterize the MLE. By the KKT conditions for convex optimization problems, the primal-dual optimal points are characterized by

(1) primal feasibility resulting in (5),
(2) dual feasibility resulting in (6) and (7),
(3) complimentary slackness resulting in (8),
(4) stationarity of the Lagrangian resulting in $\operatorname{tr}(S\hat{\Sigma}^{-1}) = p$, which is already implied by (8).

This completes the proof. □

Remark 2.3. Proposition 2.2 can easily be extended to provide properties for the existence of the MLE and a characterization of the MLE for Gaussian graphical models under MTP_2. In this case, let $G = (V, E)$ be an undirected graph. Then the primal problem has additional equality constraints, namely $K_{uv} = 0$ for all $uv \notin E$, and hence the inequality constraints in the dual problem are restricted to the entries in E, i.e., $\Sigma_{uv} \geq S_{uv}$ for all $uv \in E$. Note that if the MLE of Σ based on S exists in the Gaussian graphical model over G, it also exists in the Gaussian graphical model over G under MTP_2, since without the MTP_2 constraint the MLE needs to satisfy $\hat{\Sigma}_{uv} = S_{uv}$ for all $uv \in E$. In fact, as shown by Slawski and Hein [33], see also Theorem 3.5 below, the MLE under MTP_2 exists for all sample sizes $n \geq 2$ and this clearly remains so if graphical model constraints are added. □

We define the *maximum likelihood graph* (ML graph) $\hat{G}$ to be the graph determined by $\hat{K}$, i.e. $\hat{G} = G(\hat{K})$, where $\hat{K} = \hat{\Sigma}^{-1}$ is the MLE of K under MTP_2. We then have the following important corollary of Proposition 2.2.

Corollary 2.4. *Consider the Gaussian graphical model determined by $K_{uv} = 0$ for $uv \notin E(\hat{G})$, where $\hat{G}$ is the ML graph under* MTP_2. *Let $\bar{\Sigma}$ be the MLE of Σ under that Gaussian graphical model (without the* MTP_2 *constraint). Then $\hat{\Sigma} = \bar{\Sigma}$.*

Proof. The MLE of Σ under the Gaussian graphical model with graph $\hat{G}$ is the unique element $\bar{\Sigma} \succ 0$ that satisfies the following system of equations:

$$\begin{aligned} \bar{\Sigma}_{vv} - S_{vv} &= 0 \quad \text{for all } v \in V, \\ \bar{\Sigma}_{uv} - S_{uv} &= 0 \quad \text{for all } uv \in E(\hat{G}), \\ (\bar{\Sigma}^{-1})_{uv} &= 0 \quad \text{for all } uv \notin E(\hat{G}). \end{aligned}$$

Proposition 2.2 says that also $\hat{\Sigma}$ satisfies these equations and hence we must have $\bar{\Sigma} = \hat{\Sigma}$. □

Note that this corollary highlights the role of the complementary slackness condition (8) in inducing sparsity of the MTP_2 solution.

We emphasize that the MLE under MTP_2 is equivariant w.r.t. changes of scale so that without loss of generality we can assume that the sample covariance is normalized,

i.e. $S_{ii} = 1$ or, equivalently, $S = R$, where R is the correlation matrix. For certain of the subsequent developments this represents a convenient simplification.

Lemma 2.5. *Let S be the sample covariance matrix, R the corresponding sample correlation matrix. Denote by $\hat{\Sigma}^S$ and $\hat{\Sigma}^R$ the MLE in Proposition 2.2 based on S and R, respectively. Then*

$$\hat{\Sigma}^S_{ij} \;=\; \sqrt{S_{ii}S_{jj}}\,\hat{\Sigma}^R_{ij} \qquad \text{for all } i,j \in V.$$

Proof. Denote by D a diagonal matrix such that $D_{ii} = \sqrt{S_{ii}}$, so that $S = DRD$. The likelihood function based on S is

$$\log\det K - \operatorname{tr}(SK) \;=\; \log\det K - \operatorname{tr}(RDKD).$$

If $K' = DKD$, this can be rewritten as $\log\det K' - \operatorname{tr}(RK') - \sum_i \log S_{ii}$. Therefore, if $\hat{K}^R$ is the maximizer of $\log\det K - \operatorname{tr}(RK)$ under the MTP_2 constraints, then $D^{-1}\hat{K}^R D^{-1}$ is also an M-matrix and the maximizer of $\log\det K - \operatorname{tr}(SK)$. □

We end this section by providing simple coordinate descent algorithms for ML estimation under MTP_2. Although interior point methods run in polynomial time, for very large Gaussian graphical models it is usually more practical to apply coordinate descent algorithms. In Algorithms 1 and 2 we describe two methods for computing the MLE that only use optimization problems of size 2×2 which have a simple and explicit solution, and iteratively update the entries of K, respectively of Σ. Algorithms 1 and 2 are inspired by the corresponding algorithms for Gaussian graphical models; see, for example, [10, 35, 37]. Slawski and Hein [33] also provide a coordinate descent algorithm for estimating covariance matrices under MTP_2. However, their method updates one column/row of Σ at a time.

We first analyze Algorithm 1. Let $A = \{u, v\}$ and $B = V \setminus A$. Then note that the objective function can be written in terms of the 2×2 Schur complement $K' = K_{AA} - K_{AB}K_{BB}^{-1}K_{BA}$, since up to an additive constant

$$\log\det K - \operatorname{trace}(KS) = \log\det K' - \operatorname{trace}(K'S_{AA}).$$

Algorithm 1 Coordinate descent on K.

Input: Sample covariance matrix S, and precision ϵ.
Output: MLE $\hat{K} \in \mathcal{M}$.

1. Let $K^0 := K^1 := (\operatorname{diag}(S))^{-1}$.
2. Cycle through entries $u \neq v$ and solve the following optimization problem:

$$\begin{aligned} \underset{K \succeq 0}{\text{maximize}} \quad & \log\det(K) - \operatorname{trace}(KS) \\ \text{subject to} \quad & K_{uv} \leq 0, \\ & K_{ij} = K^1_{ij} \;\text{ for all } ij \in (V \times V) \setminus \{uu, vv, uv\}, \end{aligned}$$

   and update $K^1 = K$.
3. If $\|K^0 - K^1\|_1 < \epsilon$, set $\hat{K} = K^1$. Otherwise, set $K^0 = K^1$ and return to 2.

Algorithm 2 Coordinate descent on Σ.

Input: Sample covariance matrix $S \succ 0$, and precision ϵ.
Output: MLE $\hat{\Sigma}$ with $\hat{\Sigma}^{-1} \in \mathcal{M}$.

1. Let $\Sigma^0 := \Sigma^1 := S$
2. Cycle through entries $u \neq v$ and solve the following optimization problem:
$$\begin{aligned} \underset{\Sigma \succeq 0}{\text{maximize}} \quad & \log\det(\Sigma) \\ \text{subject to} \quad & \Sigma_{uv} \geq S_{uv}, \\ & \Sigma_{ij} = \Sigma^1_{ij} \ \text{ for all } ij \in (V \times V) \setminus \{uv\}. \end{aligned}$$
and update $\Sigma^1 = \Sigma$.
3. If $\|\Sigma^0 - \Sigma^1\|_1 < \epsilon$, set $\hat{\Sigma} = \Sigma^1$. Otherwise, set $\Sigma^0 = \Sigma^1$ and return to 2.

Defining $L := K_{AB}K_{BB}^{-1}K_{BA}$, then the optimization problem in step (2) of Algorithm 1 is equivalent to
$$\begin{aligned} \underset{K' \succeq 0}{\text{maximize}} \quad & \log\det(K') - \operatorname{trace}(K'S_{AA}) \\ \text{subject to} \quad & K'_{12} + L_{12} \leq 0. \end{aligned}$$
The unconstrained optimum to this problem is given by $K' = S_{AA}^{-1}$ and is attained if and only if $(S_{AA}^{-1})_{12} + L_{12} \leq 0$, or equivalently, if and only if
$$L_{12} \leq \frac{S_{uv}}{S_{uu}S_{vv} - S_{uv}^2}.$$
Otherwise the KKT conditions give that $K'_{12} = -L_{12}$. Maximizing over the remaining two entries of K' leads to a quadratic equation, which has one feasible solution
$$K'_{11} = \frac{1 + \sqrt{1 + 4S_{uu}S_{vv}L_{12}^2}}{2S_{uu}}, \quad K'_{22} = \frac{1 + \sqrt{1 + 4S_{uu}S_{vv}L_{12}^2}}{2S_{vv}}, \quad K'_{12} = -L_{12}. \tag{10}$$
Then the solution to the optimization problem in step (2) is given by $K_{AA} = K' + L$.

Dual to this algorithm, one can define an algorithm that iteratively updates the off-diagonal entries of Σ by maximizing the log-likelihood in direction Σ_{uv} and keeping all other entries fixed. This procedure is shown in Algorithm 2. Similarly as for Algorithm 1, the solution to the optimization problem in step (2) can be given in closed-form. Defining $A = \{u, v\}$, $B = V \setminus A$ and $L = \Sigma_{AB}\Sigma_{BB}^{-1}\Sigma_{BA}$, then analogously as in the derivation above, one can show that the solution to the optimization problem in step (2) of Algorithm 2 is given by
$$\Sigma_{uv} = \max\{S_{uv}, L_{12}\}. \tag{11}$$

We end by proving that Algorithms 1 and 2 indeed converge to the MLE. We here assume that $n \geq 2$ to guarantee existence of the MLE. Note that the suggested starting points for both algorithms can be modified.

Proposition 2.6. *Algorithms 1 and 2 converge to the MLE $\hat{K} = \hat{\Sigma}^{-1} \in \mathcal{M}$.*

Proof. The convergence to the MLE is immediate for Algorithm 2 because it is a coordinate descent method applied to a smooth and strictly concave function; see, e.g., [25]. For Algorithm 1 we use the fact that it is an example of iterative partial maximization. To prove convergence to the MLE we we will show that the assumptions of Proposition A.3 in [23] hold. The log-likelihood function that we are trying to maximize is strictly concave and so the maximum is unique. Clearly K is the maximum if and only if it is a fixed point of each update. It only remains to show that updates depend continuously on the previous value. For a given S fix K and consider a sequence of points K_n converging to K. Denote by $\tilde{K}$ and $\tilde{K}_n$ the corresponding one-step updates. We want to show that $\tilde{K}_n$ also converges to $\tilde{K}$. As above, let $A = \{u, v\}$, $B = V \setminus A$, $K' = K_{AA} - K_{AB}K_{BB}^{-1}K_{BA}$ and $L = K_{AB}K_{BB}^{-1}K_{BA}$. Outside of the block $\tilde{K}_{AA}$ this convergence is trivial; so we focus only on the three entries in $\tilde{K}_{AA}$. The function $L_{12} \mapsto (K'_{11}, K'_{22}, K'_{12})$ is continuous if and only if each coordinate is. It is clear that these functions are continuous if $L_{12} \neq \frac{S_{uv}}{S_{uu}S_{vv}-S_{uv}^2}$. It remains to show that if $L_{12} = \frac{S_{uv}}{S_{uu}S_{vv}-S_{uv}^2}$ the update in (10) gives $K' = S_{AA}^{-1}$, which can be easily checked directly. □

3. Ultrametric matrices and inverse M-matrices

In this section we exploit the link to ultrametrics in order to construct an explicit primal and dual feasible point of the maximum likelihood estimation problem. A non-negative symmetric matrix U is said to be *ultrametric* if

(i) $U_{ii} \geq U_{ij}$ for all $i, j \in \{1, \ldots, p\}$,
(ii) $U_{ij} \geq \min\{U_{ik}, U_{jk}\}$ for all $1 \leq i, j, k \leq p$.

The connection to M-matrices is established by the following result; see [9, Theorem 3.5].

Theorem 3.1. *Let U be an ultrametric matrix with strictly positive entries on the diagonal. Then U is nonsingular if and only if no two rows are equal. Moreover, if U is nonsingular then U^{-1} is an M-matrix.*

The main reason why ultrametric matrices are relevant in this paper is the following construction, which is similar to constructions used in in phylogenetics [31, Section 7.2] and single linkage clustering [15].

Let R be a symmetric $p \times p$ positive semidefinite matrix such that $R_{ii} = 1$ for all $i = 1, \ldots, p$. Consider the graph G^+ over V with an edge between i and j whenever R_{ij} is positive. Assign to each edge the corresponding positive weight R_{ij} and note that G^+ in general does not have to be connected. Define a $p \times p$ matrix Z by setting $Z_{ii} = 1$ for all i and

$$Z_{ij} := \max_{P} \min_{uv \in P} R_{uv}, \tag{12}$$

where the maximum is taken over all paths in G^+ between i and j and is set to zero if no such path exists. We call Z the *single-linkage matrix* based on R.

Example 3.2. Suppose that

$$R = \begin{bmatrix} 1 & -0.5 & 0.5 & 0.6 \\ -0.5 & 1 & 0.4 & -0.1 \\ 0.5 & 0.4 & 1 & 0.2 \\ 0.6 & -0.1 & 0.2 & 1 \end{bmatrix}$$

Then G^+ and Z are given by

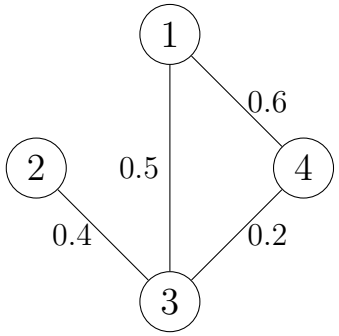

$$Z = \begin{bmatrix} 1 & 0.4 & 0.5 & 0.6 \\ 0.4 & 1 & 0.4 & 0.4 \\ 0.5 & 0.4 & 1 & 0.5 \\ 0.6 & 0.4 & 0.5 & 1 \end{bmatrix}.$$

For example, to get Z_{12} we consider two paths $1-3-2$ and $1-4-3-2$. The minimum of R_{uv} over the first path is 0.4 and over the second path 0.2. This gives $Z_{12} = 0.4$. Note that $Z \geq R$, Z is invertible, and Z^{-1} is an M-matrix. We now show that this is an example of a more general phenomenon. □

Proposition 3.3. *Let R be a symmetric $p \times p$ positive semidefinite matrix satisfying $R_{ii} = 1$ for all $i = 1, \ldots, p$. Then the single-linkage matrix Z based on R is an ultrametric matrix with $Z_{ij} \geq R_{ij}$ for all $i \neq j$. If, in addition, $R_{ij} < 1$ for all $i \neq j$, then Z is nonsingular and therefore an inverse M-matrix.*

Proof. We first show that Z is an ultrametric matrix. Z is symmetric by definition. Because R is positive semidefinite, $R_{ij} \leq 1$ for all i, j and from (12) it immediately follows that $Z_{ij} \leq 1$ and therefore $Z_{ii} \geq Z_{ij}$ for all i, j as needed. Finally, to prove condition (ii) in the definition of ultrametric, let $i, j, k \in V$. Suppose first that i, j, k lie in the same connected component of G^+. Let P_1, P_2 be the paths in G^+ such that $Z_{ik} = \min_{uv \in P_1} R_{uv}$ and $Z_{jk} = \min_{uv \in P_2} R_{uv}$. Let P_{12} be the path between i and j obtained by concatenating P_1 and P_2. Then

$$Z_{ij} = \max_{P} \min_{uv \in P} R_{uv} \geq \min_{uv \in P_{12}} R_{uv} = \min\{Z_{ik}, Z_{jk}\}.$$

Now suppose that i, j, k are not in the same connected component of G^+. In that case $0 \in \{Z_{ij}, Z_{ik}, Z_{jk}\}$. Because zero is attained at least twice, again $Z_{ij} \geq \min\{Z_{ik}, Z_{jk}\}$. Hence, Z is an ultrametric matrix. The fact that $Z_{ij} \geq R_{ij}$ for all i, j follows directly by noting that the edge ij forms a path between i and j.

Suppose now that $R_{ij} < 1$ for all $i \neq j$. In that case also $Z_{ij} < 1$ for all $i \neq j$. From this it immediately follows that no two rows of Z can be equal. Indeed, if the i-th row is equal to the j-th row for some $i \neq j$, then necessarily $Z_{ij} = Z_{ii} = Z_{jj}$, a contradiction. From Theorem 3.1 it follows that Z is an inverse M-matrix. □

As a direct consequence we obtain the following result.

Proposition 3.4. *Let S be a symmetric positive semidefinite matrix with strictly positive entries on the diagonal and such that $S_{ij} < \sqrt{S_{ii}S_{jj}}$. Then there exists an inverse M-matrix Z such that $Z \geq S$ and $Z_{ii} = S_{ii}$ for all $i = 1, \ldots, p$.*

Proof. Apply Proposition 3.3 to the normalized version R of S, with entries $R_{ij} := S_{ij}/\sqrt{S_{ii}S_{jj}}$. Because $R_{ij} < 1$ for all $i \neq j$, the corresponding single-linkage matrix Z' is ultrametric with $Z' \geq R$ and Z' is an inverse M-matrix. Define Z by $Z_{ij} = \sqrt{S_{ii}S_{jj}}Z'$. Then $Z \geq S$ and $Z_{ii} = S_{ii}$ for all $i = 1, \ldots, p$. Moreover, Z is an inverse M-matrix because Z' is. □

Proposition 3.4 is very important for our considerations. A basic application is an elegant alternative proof of the main result of [33], which says that the MLE under MTP_2 exists with probability one as long as $n \geq 2$. This is in high contrast with the existence of the MLE in Gaussian graphical models without additional constraints; see [36].

Theorem 3.5 (Slawski and Hein [33])**.** *Consider a Gaussian MTP_2 model and let S be the sample covariance matrix. If $S_{ij} < \sqrt{S_{ii}S_{jj}}$ for all $i \neq j$ then the MLE $\hat{\Sigma}$ (and $\hat{K}$) exists and it is unique. In particular, if the number n of observations satisfies $n \geq 2$, then the MLE exists with probability 1.*

Proof. The sample covariance matrix is a positive semidefinite matrix with strictly positive diagonal entries. We can apply Proposition 3.4 to obtain an inverse M-matrix Z that satisfies $Z \geq S$ and $Z_{ii} = S_{ii}$ for all i. It follows that Z satisfies primal feasibility (5) and dual feasibility (6) and (7). By Proposition 2.2 the MLE exists and it is unique by convexity of the problem. □

The matrix Z can be computed efficiently[2]. To see that, note first that in Example 3.2 we could first consider the chain T of the form $2 - 3 - 1 - 4$, which is the maximal weight spanning forest of G^+ and then construct Z by

$$Z_{ij} \;=\; \min_{uv=\overline{ij}} R_{uv},$$

where $\overline{ij}$ denotes the unique path between i and j in T. For example $Z_{12} = 0.4$, which corresponds to the minimal weight on the path $2-3-1$. This is a general phenomenon.

Suppose again that R is a symmetric $p \times p$ positive semidefinite matrix satisfying $R_{ii} = 1$ for all $i = 1, \ldots, p$. Let $\mathrm{MWSF}(R)$ be the set of all minimal maximum weight spanning forests of R. Note that all edge weights of any such forest $F \in \mathrm{MWSF}(R)$ must be positive so we must have $F \subseteq G^+$. Also, if R is an empirical correlation matrix, then $\mathrm{MWSF}(R)$ will be a singleton with probability one and in such cases we shall mostly speak of *the* MWSF.

Proposition 3.6. *The single-linkage matrix Z as defined in (12) is block diagonal with blocks corresponding to the connected components of any $F \in \mathrm{MWSF}(R)$. Within each block all elements are strictly positive and given by*

$$Z_{ij} \;=\; \min_{uv\in\overline{ij}} R_{uv},$$

[2]In our computations we use the single-linkage clustering method in R.

where $\overline{ij}$ is the unique path between i and j in a maximal weight spanning tree of R. In particular, $Z_{ij} = R_{ij}$ for all edges of $\mathrm{MWSF}(R)$.

Proof. First suppose that $i, j \in V$ lie in two different components of $F \in \mathrm{MWSF}(R)$. This means that there is no path between i and j in G^+ and so, by definition, $Z_{ij} = 0$. Because $Z_{ij} > 0$ if i, j lie in the same component of F, Z is block diagonal with blocks corresponding to connected components of $\mathrm{MWSF}(R)$.

The rest of the proof is an adaptation of a proof of a related result [31, Proposition 7.2.10]. Suppose that $i, j \in V$ lie in the same connected component of F and denote the tree in F corresponding to this component by T. By definition $Z_{ij} \geq \min_{uv \in \overline{ij}} R_{uv}$. Suppose that $Z_{ij} > \min_{uv \in \overline{ij}} R_{uv}$. We obtain the contradiction by showing that under this assumption T cannot be a maximum weight spanning tree of the corresponding connected component of G^+. Let kl be a minimum weight edge in the unique path between i and j in T. Since $Z_{ij} > R_{kl}$, there exists a path P in G^+ between i and j such that $R_{uv} > R_{kl}$ for every uv in P. Now deleting kl from T partitions the corresponding connected component of G^+ into two sets with i being in one and j being in the other block. Since P connects i and j in G^+, there must be an edge $k'l'$ (distinct from kl) in P whose end vertices lie in different blocks of this partition. Let T' be the spanning tree obtained from T by deleting kl and adding $k'l'$. Since $R_{k'l'} > R_{kl}$, the total weight of T' is greater than T, which is a contradiction. We conclude that $Z_{ij} = \min_{uv \in \overline{ij}} R_{uv}$ for all i, j in the same connected component of G^+. □

To conclude this section, we note that the starting point Σ^0 of Algorithm 2 is arbitrary as long as $\Sigma^0 \succ 0$ and $\Sigma^0 \geq S$. The single-linkage matrix constitutes another generic choice when $S = R$ is used as input. This is a particularly desirable starting point, since it can also be used when $p > n$, in which case $R \notin \mathbb{S}_{\succ 0}$ and hence not feasible as a starting point.

4. The maximum likelihood graph

Fitting a Gaussian model with MTP_2 constraints tends to induce sparsity in the maximum likelihood estimate $\hat{K}$. In this section, we analyze the sparsity pattern that arises in this way. We assume again without loss of generality that $S = R$ is a sample correlation matrix so that $R_{ii} = 1$ for all i and $R_{ij} < 1$ for all $i \neq j$. Consider again the weighted graph G^+ on p nodes with edges corresponding precisely to the positive off-diagonal entries of R with weights R_{ij}. We begin this section with a basic lemma that reduces our analysis to the case where the graph G^+ is connected.

Lemma 4.1. *The MLE $\hat{\Sigma}$ under* MTP_2 *is a block diagonal matrix with strictly positive entries in each block. The blocks correspond precisely to trees in* $\mathrm{MWSF}(R)$.

Proof. Firstly, since $\hat{\Sigma}$ is an inverse M-matrix, it is block diagonal with strictly positive entries in each block; see, e.g., Theorem 4.8 in [19]. We will show that each block of $\hat{\Sigma}$ corresponds precisely to a tree in $\mathrm{MWSF}(R)$.

Denote the vertex sets for a forest $F \in \mathrm{MWSF}(R)$ as $T_1, \ldots, T_k$ and the blocks of $\hat{\Sigma}$ as $B_1, \ldots, B_l$. Firstly, for any T_i there must be a j so that $T_i \subseteq B_j$; this is true since

all entries in R along the edges of T_i are positive and thus $\hat{\Sigma}_{uv} \geq R_{uv} > 0$. Thus the block partitioning corresponding to the trees is necessarily finer than that of $\hat{\Sigma}$.

On the other hand, suppose that two different trees T_i and T_j in F are in the same block of $\hat{\Sigma}$ so that $\hat{\Sigma}_{uv} > 0$ for all $u \in T_i$ and $v \in T_j$. Then, as we must have $R_{uv} \leq 0$, also necessarily $\hat{\Sigma}_{uv} - R_{uv} > 0$. Complementary slackness (8) now implies that $\hat{K}_{uv} = 0$ for all $u \in T_i$ and $v \in T_j$, and hence $\hat{K}$ is block-diagonal with blocks corresponding to the trees in F. Since $\hat{\Sigma} = \hat{K}^{-1}$, we also get $\hat{\Sigma}_{uv} = 0$ which contradicts that u and v are in the same block of $\hat{\Sigma}$. □

This result shows that, without loss of generality, we can always assume that G^+ is connected and then $\mathrm{MWSF}(R) = \mathrm{MWST}(R)$ consists of trees only. If there are more than one connected component, we simply compute the MLE for each component separately and combine them together in block diagonal form. Hence, from now on we always assume that all forests in $\mathrm{MWSF}(R)$ are just trees.

The following two theorems are the main results of the present paper.

Theorem 4.2. *Let $\hat{\Sigma}$ be the MLE of the correlation matrix under* MTP_2 *based on the sample correlation matrix R. Then any $F \in \mathrm{MWSF}(R)$ is a subgraph of the ML graph $\hat{G} = G(\hat{K})$ with probability one.*

Proof. Let $\bar{\Sigma}$ denote a solution to the following optimization problem:

$$\begin{aligned} \underset{\Sigma}{\text{maximize}} \quad & \log\det(\Sigma) \\ \text{subject to} \quad & \Sigma_{ii} = 1, \quad \text{for all } i = 1, \ldots, p, \\ & \Sigma_{ij} \geq R_{ij}, \quad \text{for all } i \neq j, \\ & \Sigma_{ij} = R_{ij}, \quad \text{for all } ij \in \mathrm{MWSF}(R), \\ & \Sigma \text{ is an inverse M-matrix}, \\ & \Sigma \succeq 0. \end{aligned}$$

Note that $\bar{\Sigma}$ exists since the single linkage matrix Z is a feasible solution according to Proposition 3.3 and Proposition 3.6, and the objective function is bounded because the maximum is attained without the constraint that Σ is an inverse M-matrix.

Next we want to show that $\hat{\Sigma} = \bar{\Sigma}$. We prove this by applying a step of Algorithm 2 on an entry $kl \in F \in \mathrm{MWSF}(R)$. Without loss of generality we assume that $kl = 12 \in F$ and fix $B = V \setminus \{1, 2\}$. Because $\bar{\Sigma}$ is an inverse M-matrix, $\bar{\Sigma}_{12} \geq \bar{\Sigma}_{1B}(\bar{\Sigma}_{BB})^{-1}\bar{\Sigma}_{B2} = L_{12}$. Also $\bar{\Sigma}_{12} = R_{12}$ by construction and so $R_{12} \geq L_{12}$. It follows that the update (11) makes no change in $\bar{\Sigma}$, which proves that $\hat{\Sigma} = \bar{\Sigma}$.

The last argument shows that $\hat{\Sigma}_{ij} = R_{ij}$ for any edge in a maximal weight spanning forest of R. In other words, by Corollary 2.4, $\hat{\Sigma} = \bar{\Sigma}$ is also the MLE under the Gaussian graphical model with graph $\bar{G} = \mathrm{MWSF}(R) \cup G(\hat{K})$. However, this happens with probability zero unless $\mathrm{MWSF}(R) \subseteq G(\hat{K})$ and hence the conclusion follows. □

At the heart of the proof of Theorem 4.2 we showed the following remarkable property of the dual convex optimization problems (2) and (9): *dual constraints* (i.e., $\Sigma_{ij} \geq R_{ij}$)

that are active at an initial point that is primal feasible (i.e., Σ is an inverse M-matrix) *stay active at the solution*. The single linkage matrix Z is an example of an inverse M-matrix satisfying $Z_{ij} = R_{ij}$ for all $ij \in \mathrm{MWSF}(R)$ and hence $\mathrm{MWSF}(R)$ is a subgraph of $\hat{G}$. As a consequence, the statement of Theorem 4.2 can be strengthened: given any inverse M-matrix Z' such that $Z' \geq R$ and $Z'_{ii} = R_{ii}$ for all i, then $ij \in \hat{G}$ for all $i \neq j$ such that $Z'_{ij} = R_{ij}$.

Theorem 4.2 gives a partial justification for the heuristics used in applied statistics to learn high-dimensional Gaussian graphical models by initializing search procedures in $\mathrm{MWSF}(|R|)$, see for example [17]. We emphasize that the result is not true without the MTP_2 assumption.

An algebraic analogue of Theorem 4.2 holds for general inverse M-matrices. More precisely, we have the following result:

Theorem 4.3. *Assume that C is an inverse M-matrix with $C_{ii} = 1$ and let $K = C^{-1}$. Then if $\mathrm{MWSF}(C)$ consists of a single element, it holds that $\mathrm{MWSF}(C) \subseteq G(K)$.*

Proof. Let A denote the adjacency matrix of the complement of $G(K)$ and let $R = C - \varepsilon A$. If ε is sufficiently small, we have $\mathrm{MWSF}(R) = \mathrm{MWSF}(C)$ and $\hat{\Sigma}(R) = C$ since C fulfills all the conditions in Proposition 2.2. Now we proceed as in the proof of Theorem 4.2 and note that for $uv \notin G(K)$, we have $C_{uv} - R_{uv} = \varepsilon > 0$ which by complementary slackness implies $K_{uv} = 0$ and hence all edges in $\mathrm{MWST}(R) = \mathrm{MWSF}(C)$ must be in $G(K)$. □

Note that Theorem 4.2 and Theorem 4.3 represent a very strong form of absence of a paradox of Yule–Simpson type for Gaussian MTP_2 distributions. For MTP_2 distributions, it clearly holds that a positive correlation cannot change sign when conditioning with other variables. The results in these theorems are stronger and say that for Gaussian MTP_2 distributions, correlations along edges of $\mathrm{MWSF}(C)$ necessarily correspond to direct associations in the concentration graph and are not resulting from indirect associations.

We have seen that the $\mathrm{MWSF}(R)$ provides a lower bound for the ML graph $\hat{G} = G(\hat{K})$. Below we provide a simple procedure for identifying an upper bound for $\hat{G}$. This procedure relies on the estimation of the standard Gaussian graphical model over the tree $\mathrm{MWSF}(R)$. The MLE under this assumption, denoted by $\tilde{\Sigma}$, can be computed very efficiently and it satisfies

$$\tilde{\Sigma}_{ij} \;=\; \prod_{(u,v)\in\overline{ij}} R_{uv}.$$

where $\overline{ij}$ denotes the unique path between i and j in $\mathrm{MWSF}(R)$; see, for example, [38, Section 8.2].

Proposition 4.4. *The pair ij forms an edge in the ML graph only if*

$$R_{ij} \;\geq\; \tilde{\Sigma}_{ij} \;=\; \prod_{(u,v)\in\overline{ij}} R_{uv}.$$

In particular $R_{ij} \leq 0$ implies that ij is not an edge of the ML graph.

Proof. Because $\hat{\Sigma}$ is an inverse M-matrix it is necessarily a path product matrix; see [18, Section 1]. In particular, for all i, j

$$\hat{\Sigma}_{ij} \;\geq \prod_{(u,v)\in\overline{ij}} \hat{\Sigma}_{uv}.$$

By Corollary 2.4, the MLE under MTP_2 is equal to the MLE under the Gaussian graphical model over the ML graph. Therefore, if ij forms an edge in the ML graph, then $\hat{\Sigma}_{ij} = R_{ij}$. By Theorem 4.2, all edges of $\mathrm{MWSF}(R)$ lie in the ML graph and so $\hat{\Sigma}_{uv} = R_{uv}$ for all edges of $\overline{ij}$. This proves that the inequality in the statement of the proposition must hold. □

Motivated by this result we define the *excess correlation graph* (EC graph) $\mathrm{EC}(R)$ of R by the condition

$$i \sim j \iff R_{ij} \geq \prod_{(u,v)\in\overline{ij}} R_{uv}.$$

Thus the EC graph has edges ij whenever the observed correlation between i and j is in excess of or equal to what is explained by the spanning forest; by construction, we have with probability one that

$$\mathrm{MWSF}(R) \subseteq \hat{G} = G(\hat{K}) \subseteq \mathrm{EC}(R).$$

The inclusions above are typically strict. For example, if R is an inverse M-matrix, then $\mathrm{EC}(R)$ is the complete graph, whereas $G(\hat{K})$ can be arbitrary; this follows from [12, Proposition 6.3].

The following construction may be useful for studying MTP_2 Gaussian distributions. Let R be the sample correlation matrix. Define the matrix W by

$$W_{ij} \;:=\; \max_{P} \prod_{(u,v)\in P} R_{uv},$$

where, like in (12), the maximum is taken over all paths in G^+ between i and j and is set to zero if no such path exists. Transforming $D_{ij} = -\log R_{ij}$ gives a distance based interpretation, in which W_{ij} is related to the shortest distance between i and j in G^+ with edge lengths given by D_{uv}.

It is easy to show that $Z \geq W \geq R$ and that W is always equal to the MLE $\hat{\Sigma}$ in the case when $p \leq 3$. For general p we do not know conditions on R that assure that W is an inverse M-matrix, or the MLE. We note that, quite surprisingly, $W = \hat{\Sigma}$ for the `carcass` data discussed in the introduction.

To complete the discussion of the ML graph, we provide a more detailed analysis in the case when $n = 2$.

Proposition 4.5. *Let R be a $p \times p$ sample correlation matrix of rank 2 such that $R_{ij} < 1$ for all $i \neq j$. Then the ML graph $G(\hat{K})$ consists of at most two paths and thus $\mathrm{MWSF}(R) = G(\hat{K})$.*

Proof. R can be decomposed into $R = XX^T$, where X is a matrix of size $p \times 2$ consisting of row vectors of length 1. The p row vectors of R are denoted by $x_1, \ldots, x_p$. Without loss of generality we can assume that all vectors lie in the upper half-plane and that $x_1 = (1, 0)$. In addition, without loss of generality we assume that the vectors $x_1, \ldots, x_p$ are arranged as shown in Figure 2. We denote by θ_{ij} the angle between x_i and x_j. Note that $R_{ij} = \cos(\theta_{ij})$. Since $\cos(x)$ is decreasing on $[0, \pi]$, the maximum weight spanning tree is the path $(1, 2, 3, \ldots, p)$, assuming that all angles $\theta_{i,i+1} \leq \pi/2$ (otherwise the maximum weight spanning forest has two paths). By the previous theorem this implies that these paths are contained in the ML graph. Buhl proved in [6] that the MLE for the cycle with 2 observations exists if and only if the vectors $x_1, \ldots, x_p$ do not follow the same ordering as the cycle ordering. This implies that the MLE cannot be a cycle. The result follows. □

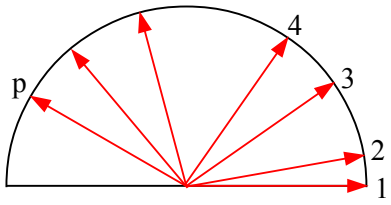

FIGURE 2. Configuration of vectors $x_1, \ldots, x_p$.

Note that since the ML graph is a forest in this case, the EC graph will consist of complete graphs corresponding to the (at most two) connected components of the ML graph.

5. Gaussian signed MTP_2 distributions

In this section we discuss how our results can be generalized to so-called signed MTP_2 Gaussian distributions, where the distribution is MTP_2 up to sign swapping. Such distributions were discussed by Karlin and Rinott [21]. More precisely, a random variable X has a *signed* MTP_2 *distribution* if there exists a diagonal matrix D with $D_{ii} = \pm 1$ such that DX is MTP_2. In the context of Gaussian graphical models, signed MTP_2 distributions were studied in [26], where they are called *non-frustrated*.

The following characterization of signed MTP_2 Gaussian distributions is a direct consequence of [21, Theorem 3.1 and Remark 1.3].

Proposition 5.1. *A Gaussian random variable X has a signed* MTP_2 *distribution if and only if $|X|$ is* MTP_2.

Our first result in this section is an alternative characterization of signed MTP_2 Gaussian distributions via inequalities on the entries of the concentration matrix. The proof is somewhat technical and can be found in the appendix.

Theorem 5.2. *A Gaussian random variable X with concentration matrix K has a signed* MTP_2 *distribution if and only if for every cycle $(i_1, \ldots, i_k, i_1)$ in the graph $G(K)$*

$$(-1)^k K_{i_1 i_2} K_{i_2 i_3} \cdots K_{i_k i_1} > 0. \tag{13}$$

Signed MTP_2 distributions are relevant for example because of their appearance when studying tree models.

Proposition 5.3. *Every Gaussian graphical model over a tree consists of signed* MTP_2 *distributions. These* MTP_2 *distributions are precisely the ones in which the covariance matrix* Σ *has only nonnegative entries.*

Proof. Let T be a tree and $K = \Sigma^{-1}$ be a concentration matrix in the Gaussian graphical model over T. Then $G(K)$ is a subgraph of T and in particular it has no cycles. Hence by Theorem 5.2 it is signed MTP_2. The second part of the statement follows from [12, Corollary 7.3]. □

Because signed MTP_2 distributions are closed under taking margins, Proposition 5.3 can be further generalized. The following theorem covers, in particular, Examples 4.1-4.5 in [21].

Theorem 5.4. *Every distribution on a Gaussian tree model with hidden variables is signed* MTP_2.

Gaussian tree models with hidden variables have many applications, in partiular related to modeling evolutionary processes; see, e.g., [7, 32]. As an important submodel they contain the Brownian motion tree model [13]. Another example of a Gaussian tree model is the factor analysis model with a single factor; it corresponds to a Gaussian model on a star tree, whose inner node is hidden. The MTP_2 distributions in this model correspond to the distributions in a Spearman model [24, 34], where the hidden factor is interpreted as intelligence.

Let R be a sample correlation matrix. Maximizing the likelihood over all signed MTP_2 Gaussian distributions corresponds to determining the sign matrix D, with $D_{ii} = \pm 1$, that maximizes the likelihood for all 2^p possible matrices DRD. A natural heuristic is to choose D^* such that $D^*_{ii}D^*_{jj}R_{ij} \geq 0$ for all edges ij of $\mathrm{MWSF}(|R|)$. We provide conditions under which this procedure indeed gives the MLE under signed MTP_2, and we also provide examples showing that this is not true in general.

First we describe how to obtain a sign swapping matrix D^* such that $D^*_{ii}D^*_{jj}R_{ij} \geq 0$ for all edges ij of $\mathrm{MWSF}(|R|)$. Root $\mathrm{MWSF}(|R|)$ at node 1, that is, regard $\mathrm{MWSF}(|R|)$ as a directed tree with all edges directed away from 1. Set $D^*_{11} = 1$. Then proceed recursively. For any edge $i \to j$ suppose that D^*_{ii} is known and set $D^*_{jj} := \mathrm{sgn}(D^*_{ii}R_{ij})$. Note that by construction

$$D^*_{ii} := \mathrm{sgn}(R_{1i_1}R_{i_1i_2}\cdots R_{i_ki}), \tag{14}$$

where $1 \to i_1 \to i_2 \to \cdots \to i_k \to i$ is the unique path from 1 to i in $\mathrm{MWSF}(|R|)$. We set $D^*_{ii} = 0$ if no such path exists. It is easy to check that the resulting D^* satisfies $D^*_{ii}D^*_{jj}R_{ij} \geq 0$ for all edges ij of $\mathrm{MWSF}(|R|)$.

Proposition 5.5. *Suppose that* R *is a sample correlation matrix such that*

$$R_{ij}R_{ik}R_{jk} > 0 \quad \textit{for all } i,j,k \in V. \tag{15}$$

Then the MLE based on R *over signed* MTP_2 *Gaussian distributions is equal to the MLE based on the sample correlation matrix* D^*RD^* *over* MTP_2 *distributions.*

Proof. Let $i_1 - \ldots - i_k - i_1$ be an undirected cycle. Then (15) implies that

$$R_{i_1 i_2} R_{i_2 i_3} \cdots R_{i_k i_1} > 0, \tag{16}$$

because the sign of this expression is equal to the sign of

$$(R_{i_1 i_2} R_{i_2 i_3} R_{i_1 i_3})(R_{i_1 i_3} R_{i_3 i_4} R_{i_1 i_4})(R_{i_1 i_4} R_{i_4 i_5} R_{i_1 i_5}) \cdots (R_{i_1 i_k} R_{i_k i_1} R_{i_1 i_1}).$$

We now use this to show that $D^* R D^*$ has only positive entries. Let i, j be any two nodes and let $1 \to i_1 \to \cdots \to i_k \to i$ and $1 \to j_1 \to \cdots \to j_l \to j$ be the paths from 1 to i and j, respectively. By (14) we obtain

$$\operatorname{sgn}(D^*_{ii} D^*_{jj} R_{ij}) \;=\; \operatorname{sgn}(R_{1 i_1} \cdots R_{i_k i} R_{ij} R_{j j_l} \cdots R_{j_1 1}),$$

which is positive by (16).

So without loss of generality we can assume that all entries of R are nonnegative and hence that D^* is the identity matrix $\mathbb{I}_p$. To complete the proof, we need to show that the likelihood over MTP_2 distributions given the sample correlation matrix DRD is maximized by $D = \mathbb{I}_p$. This is because $(D_{ii} D_{jj} - 1) \leq 0$ and $R_{ij} K_{ij} \leq 0$, and hence

$$\ell(K; R) - \ell(K; DRD) = \operatorname{tr}(DRDK) - \operatorname{tr}(RK) = \sum_{i,j} (D_{ii} D_{jj} - 1) R_{ij} K_{ij} \geq 0,$$

which completes the proof. □

By combining Proposition 5.5 with Proposition 4.5, we obtain the following result.

Corollary 5.6. *If the sample correlation matrix R is based on $n = 2$ observations, then the MLE over signed* MTP_2 *Gaussian distributions given R is equal to the MLE over* MTP_2 *Gaussian distributions given the modified sample correlation matrix $D^* R D^*$ and consists of a single path.*

Proposition 5.5 does not extend to an arbitrary sample correlation matrix. We give a simple counterexample.

Example 5.7. Suppose that the sample correlation matrix is

$$R = \begin{bmatrix} 1 & 0.3 & 0.11 & 0.3 \\ 0.3 & 1 & -0.1 & -0.1 \\ 0.11 & -0.1 & 1 & -0.1 \\ 0.3 & -0.1 & -0.1 & 1 \end{bmatrix}.$$

Then $\mathrm{MWSF}(|R|)$ is given by the star graph with edges $1-2$, $1-3$, $1-4$. Since R is positive on these entries, $D^* = \mathbb{I}_p$. But one can check that the corresponding MLE has a lower likelihood than the MLE after changing the sign of the third variable.

The intuition is the following. The log-likelihood based on R is up to an additive constant given by

$$\begin{aligned} \underset{\Sigma}{\text{minimize}} \quad & -\log\det(\Sigma) \\ \text{subject to} \quad & \Sigma_{11} = \Sigma_{22} = \Sigma_{33} = \Sigma_{44} = 1, \\ & \Sigma_{12} \geq R_{12},\ \Sigma_{13} \geq R_{13},\ \Sigma_{14} \geq R_{14}, \\ & \Sigma_{23} \geq 0,\ \Sigma_{24} \geq 0,\ \Sigma_{34} \geq 0, \\ & \Sigma \succeq 0. \end{aligned}$$

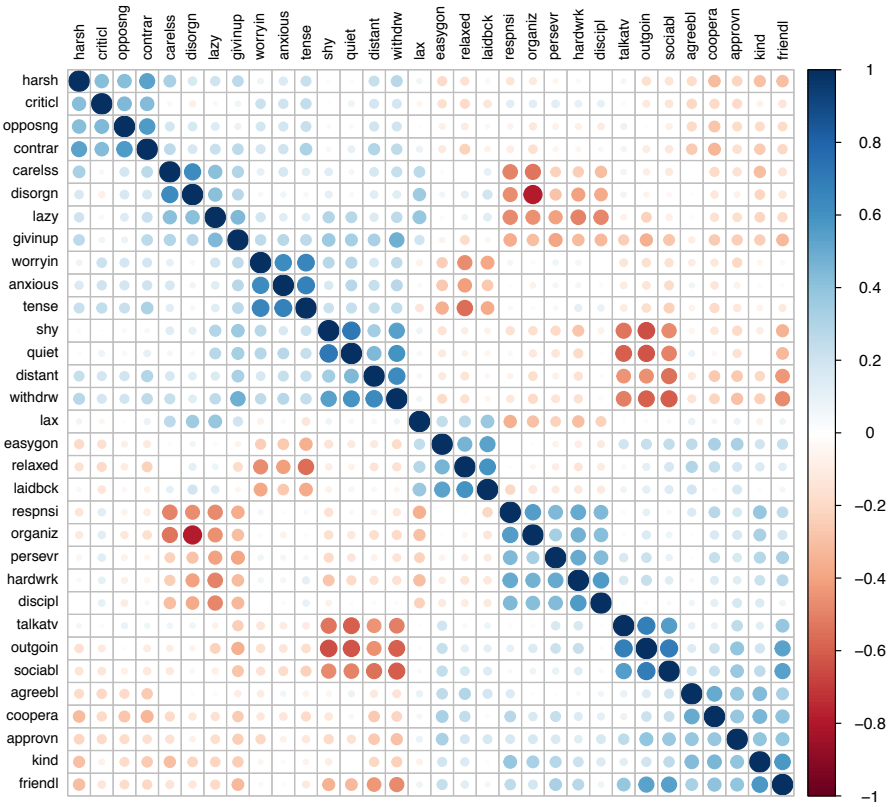

FIGURE 3. Correlation matrix of personality traits from the data set described in [27].

By changing the sign of the third variable, we replace the constraint $1-3$ by two constraints $2-3$ and $3-4$. The resulting optimization problem is

$$\begin{aligned} \underset{\Sigma}{\text{minimize}} \quad & -\log\det(\Sigma) \\ \text{subject to} \quad & \Sigma_{11}=\Sigma_{22}=\Sigma_{33}=\Sigma_{44}=1, \\ & \Sigma_{12}\geq R_{12},\ \Sigma_{14}\geq R_{14},\ \Sigma_{23}\geq -R_{23},\ \Sigma_{34}\geq -R_{34}, \\ & \Sigma_{13}\geq 0,\ \Sigma_{24}\geq 0, \\ & \Sigma\succeq 0. \end{aligned}$$

Note that R_{13} is only slightly larger than $-R_{23}$ and $-R_{24}$. Hence, in essence we are increasing the number of constraints by one, which explains the decrease of the log-likelihood value.

We end this paper by illustrating how our results can be applied to factor analysis in psychometrics.

Example 5.8. Single factor models are routinely used to study the personalities in psychometrics. Consider the following example from [27][3]: 240 individuals were asked to rate themselves on the scale 1-9 with respect to 32 different personality traits. The

[3]We downloaded the data from http://web.stanford.edu/class/psych253/tutorials/FactorAnalysis.html.

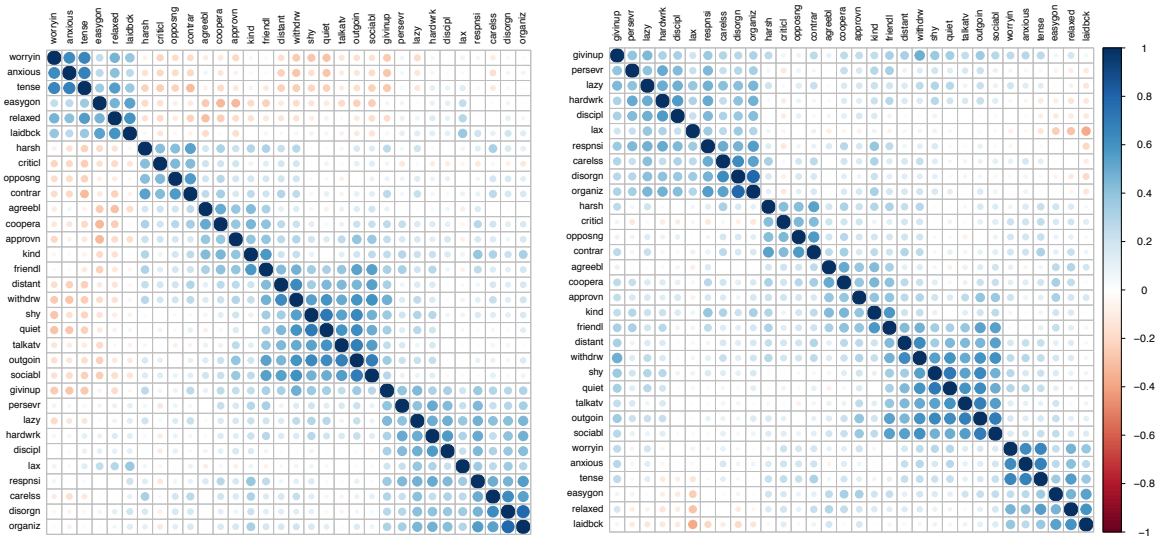

FIGURE 4. The correlation matrix of the data set on personality traits after performing the sign switches as defined in (14) is shown on the left. The correlation matrix resulting from switching the signs of the 16 (negative) traits that constitute the first block of variables in Figure 3 is shown on the right.

resulting correlation matrix is shown in Figure 3. It appears to have a block structure with predominantly positive entries in each diagonal block and negative entries in the off-diagonal block. Also analyzing the respective variables, they seem to correspond to positive and negative traits. It is therefore natural to assume that this data follows a signed MTP_2 distribution and analyze it under this constraint.

The correlation matrix resulting from the sign switching procedure described in (14) is shown on the left in Figure 4, while the correlation matrix resulting from switching the signs of the 16 (negative) traits that constitute the first block of variables in Figure 3 is shown on the right in Figure 4. These plots suggest that the matrix on the right is closer to being MTP_2. In fact, its log-likelihood (i.e., the value of $\frac{n}{2}(\log\det K - \mathrm{tr}(SK))$) is -2046.146, as compared to the log-likelihood value of -2071.717 resulting from the sign switching procedure described in (14). For comparison, the value of the unconstrained log-likelihood is -1725.075 and the value of the log-likelihood under MTP_2 without sign switching is -2356.639. The graphical models based on no sign switching and switching the signs of the 16 negative traits are shown in Figure 5. The unconstrained log-likelihood gives a lower bound of 642.142 on the likelihood ratio statistic to test signed MTP_2 constraints, while the likelihood ratio statistic to test MTP_2 constraints against the saturated model is equal to 1263.128.

APPENDIX A. PROOF OF THEOREM 5.2

We start by proving the following auxiliary result.

Proposition A.1. *A (non-degenerate) Gaussian random variable X has a signed* MTP_2 *distribution if and only if its concentration matrix lies in the closure of the*

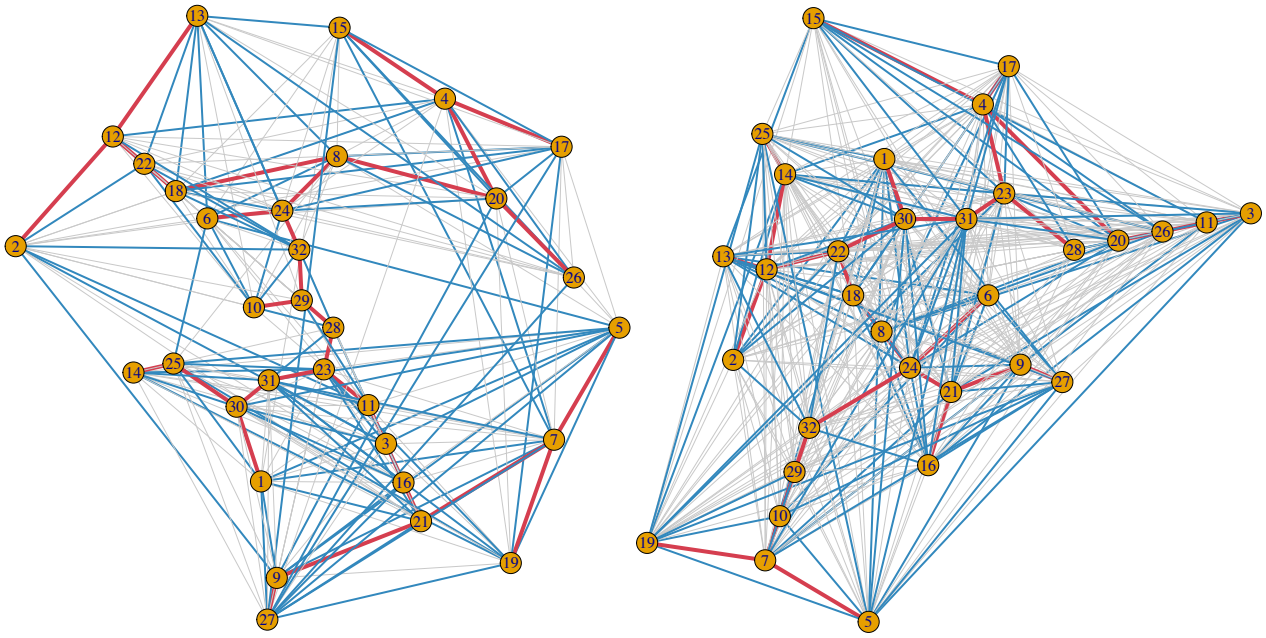

FIGURE 5. The graphical models resulting from estimation under MTP_2 based on the correlation matrix shown in Figure 3 and the correlation matrix shown in Figure 4 (right). The thin gray edges correspond to the edges of the EC graph that are not part of the ML graph. The blue edges represent edges of the ML graph that are not part of the minimum weight spanning tree. The latter is represented by thick red edges.

set satisfying

$$K_{ij}K_{ik}K_{jk} < 0 \qquad \textit{for all } 1 \le i < j < k \le p. \tag{17}$$

Proof. If $\tilde{K} := DKD$ is an M-matrix then

$$\tilde{K}_{ij}\tilde{K}_{ik}\tilde{K}_{jk} \;=\; D_{ii}^2 D_{jj}^2 D_{kk}^2 K_{ij}K_{ik}K_{jk} \;=\; K_{ij}K_{ik}K_{jk}$$

and so $K_{ij}K_{ik}K_{jk} \le 0$ for all $1 \le i < j < k \le p$. If $\tilde{K}$ has no zeros, then K lies in the set defined by (17) and so also in its closure. If $\tilde{K}$ has zeros, there exists a sequence $\tilde{K}^{(n)}$ of matrices with strictly negative off-diagonal elements such that $\tilde{K}^{(n)} \to \tilde{K}$ elementwise. In that case also $D\tilde{K}^{(n)}D \to D\tilde{K}D = K$ and so K lies in the closure of the set defined by (17). This proves the 'only if' direction.

Now suppose that K lies in the closure of the set given by (17). First, consider the case, when K satisfies (17) and thus, in particular, $K_{ij} \neq 0$ for all $i < j$. Then

$$\tilde{K}_{ij} := (\mathrm{sgn}K_{1i}\mathrm{sgn}K_{1j})K_{ij} = -(\mathrm{sgn}K_{ij})K_{ij} = -|K_{ij}| < 0,$$

where the second equation follows from the fact that $K_{1i}K_{1j}K_{ij} < 0$. In particular, $\tilde{K} = [\tilde{K}_{ij}]$ is an M-matrix. If some K_{ij} are zero, then we take a sequence $K^{(n)}$ of matrices satisfying (17) such that $K^{(n)} \to K$. For each $K^{(n)}$ we determine $D^{(n)}$ such that $D^{(n)}K^{(n)}D^{(n)}$ is an M-matrix. By continuity of matrix multiplication and the fact

that the set of M-matrices is relatively closed, there must exist D such that DKD is an M-matrix and $D = \lim D^{(n)}$. □

Remark A.2. It may look surprising at first but the closure of the set given by (17) is *not* the set with strict inequalities replaced by $\leq$. For example, consider a 4×4-matrix K such that $K_{13} = K_{24} = 0$, $K_{12} > 0$, and $K_{23}, K_{34}, K_{14} < 0$. Then for all triples it holds that $K_{ij}K_{ik}K_{jk} \leq 0$ (in fact "$= 0$"), but there is no sequence $K_{13}^{(n)} \to 0$ such that $K_{12}K_{13}^{(n)}K_{23} < 0$ and $K_{13}^{(n)}K_{14}K_{34} < 0$.

Proof of Theorem 5.2. To prove necessity, note that every point that lies in the closure of (17) satisfies also the non-strict version of the inequalities in (17). The inequalities $K_{ij}K_{ik}K_{jk} \leq 0$ for triples $\{i_1, i_2, i_3\}$, $\{i_1, i_3, i_4\}$, $\{i_1, i_4, i_5\}$, …, $\{i_1, i_{k-1}, i_k\}$ multiplied together yield $(-1)^k K_{i_1 i_2} K_{i_2 i_3} \cdots K_{i_k i_1} \geq 0$. Because all elements in this product are nonzero, (13) must hold.

To prove sufficiency, by Proposition A.1 it suffices to show that every point K satisfying (13) lies in the closure of the set described by (17). If there is a vertex, say vertex 1, in $G(K)$ that is linked by an edge to any other vertex, we can proceed as in the proof of Proposition A.1. Let D be a diagonal matrix such that $D_{ii} = -\mathrm{sgn}K_{1i}$. Let $\tilde{K} = DKD$. If $K_{ij} = 0$, then $\tilde{K}_{ij} = 0$. If $K_{ij} \neq 0$, then

$$\tilde{K}_{ij} = D_{ii}D_{jj}K_{ij} = \mathrm{sgn}(K_{1i})\mathrm{sgn}(K_{1j})K_{ij} = -|K_{ij}|,$$

where the last equality follows from the fact that necessarily $K_{1i}K_{1j}K_{ij} < 0$. Hence $\tilde{K}$ is an M-matrix and thus K is the concentration matrix of a signed MTP_2 Gaussian distribution. As a consequence, it lies in the closure of (17).

Now suppose that $G(K)$ does not contain a vertex linked by an edge to any other vertex. If $G(K)$ is not connected we can use an inductive argument because K must be block-diagonal. Hence we can assume that $G(K)$ is connected. To show that K lies in the closure of (17) we construct a sequence $K^{(n)} \to K$ as follows. Let $K_{ij}^{(n)} = K_{ij}$ whenever $i, j \neq 1$. If $K_{1j} \neq 0$ then let $K_{1j}^{(n)} = K_{1j}$ for all $n \geq 1$. If $K_{1j} = 0$ then consider any path $(1, i_1, \ldots, i_k, j)$ from 1 to j in $G(K)$. The sign $s \in \{-1, +1\}$ of the expression $(-1)^k K_{1i_1} K_{i_1 i_2} \cdots K_{i_k j}$ does not depend on the choice of the path. Indeed, if $(1, j_1, \ldots, j_l, j)$ is any other path, then the walk $(1, i_1, \ldots, i_k, j, j_l, \ldots, j_1, 1)$ can be decomposed into cycles in $G(K)$. For each of these cycles we have inequality (13). Multiplying all these inequalities together yields

$$(-1)^k K_{1i_1} K_{i_1 i_2} \cdots K_{i_k j} \cdot (-1)^l K_{j j_l} \cdots K_{j_2 j_1} \;\geq\; 0.$$

Hence the signs of the two paths are either both negative or both positive. Now for all j such that $K_{1j} = 0$ set $K_{1j}^{(n)} = -\frac{s}{n}$. By construction, $K^{(n)} \to K$ and $K_{1j}^{(n)} \neq 0$ for all $j \in V$ and $n \geq 1$. Then as in the previous paragraph we construct a diagonal matrix $D^{(n)}$ such that $D_{ii}^{(n)} = -\mathrm{sgn}(K_{1i}^{(n)})$ and let $\tilde{K}^{(n)} = D^{(n)}K^{(n)}D^{(n)}$. If $K_{ij}^{(n)} = 0$, then $\tilde{K}_{ij}^{(n)} = 0$. If $K_{ij}^{(n)} \neq 0$, then

$$\tilde{K}_{ij}^{(n)} = \mathrm{sgn}(K_{1i}^{(n)})\mathrm{sgn}(K_{1j}^{(n)})K_{ij}^{(n)} = -|K_{ij}^{(n)}|.$$

The proof of the last equality is based on the fact that $K_{1i}^{(n)}K_{1j}^{(n)}K_{ij}^{(n)} < 0$, which we now show to conclude the proof. By construction, the sign of $K_{1i}^{(n)}K_{1j}^{(n)}K_{ij}^{(n)}$ is equal to

the sign of

$$(-1)^{k+l} K_{1i_1} \cdots K_{i_k i} K_{ij} K_{jj_l} \cdots K_{i_1 1}, \tag{18}$$

where $(1, i_1, \ldots, i_k, i)$ and $(1, j_1, \ldots, j_l, j)$ are any two paths in G_K; take $(1, i)$ if $K_{1i} \neq 0$, and $(1, j)$ if $K_{1j} \neq 0$. By splitting the walk $(1, i_2, \ldots, i_k, i, j, j_l, \ldots, i_1, 1)$ into cycles, we can use (13) to show that (18) must be negative. □

References

[1] A. Anandkumar, V. Y. Tan, F. Huang, and A. S. Willsky, *High-dimensional Gaussian graphical model selection: Walk summability and local separation criterion*, The Journal of Machine Learning Research, 13 (2012), pp. 2293–2337.

[2] F. Bartolucci and J. Besag, *A recursive algorithm for Markov random fields*, Biometrika, 89 (2002), pp. 724–730.

[3] F. Bartolucci and A. Forcina, *A likelihood ratio test for* MTP_2 *within binary variables*, Ann. Statist., 28 (2000), pp. 1206–1218.

[4] B. Bhattacharya, *Covariance selection and multivariate dependence*, Journal of Multivariate Analysis, 106 (2012), pp. 212–228.

[5] E. Bølviken, *Probability inequalities for the multivariate normal with non-negative partial correlations*, Scand. J. Statist., 9 (1982), pp. 49–58.

[6] S. L. Buhl, *On the existence of maximum likelihood estimators for graphical Gaussian models*, Scand. J. Statist., 20 (1993), pp. 263–270.

[7] M. J. Choi, V. Y. F. Tan, A. Anandkumar, and A. S. Willsky, *Learning latent tree graphical models*, J. Mach. Learn. Res., 12 (2011), pp. 1771–1812.

[8] A. Colangelo, M. Scarsini, and M. Shaked, *Some notions of multivariate positive dependence*, Insurance: Mathematics and Economics, 37 (2005), pp. 13–26.

[9] C. Dellacherie, S. Martinez, and J. San Martin, *Inverse M-matrices and ultrametric matrices*, vol. 2118, Springer, 2014.

[10] A. P. Dempster, *Covariance selection*, Biometrics, (1972), pp. 157–175.

[11] J. Djolonga and A. Krause, *Scalable variational inference in log-supermodular models*, In International Conference on Machine Learning (ICML), (2015).

[12] S. Fallat, S. L. Lauritzen, K. Sadeghi, C. Uhler, N. Wermuth, and P. Zwiernik, *Total positivity in Markov structures*, to appear in the Annals of Statistics, (2016).

[13] J. Felsenstein, *Maximum-likelihood estimation of evolutionary trees from continuous characters.*, American Journal of Human Genetics, 25 (1973), pp. 471–492.

[14] C. M. Fortuin, P. W. Kasteleyn, and J. Ginibre, *Correlation inequalities on some partially ordered sets*, Comm. Math. Phys., 22 (1971), pp. 89–103.

[15] J. C. Gower and G. Ross, *Minimum spanning trees and single linkage cluster analysis*, Applied statistics, (1969), pp. 54–64.

[16] M. Grant and S. Boyd, *CVX: Matlab software for disciplined convex programming, version 2.1*. http://cvxr.com/cvx, Mar. 2014.

[17] S. Højsgaard, D. Edwards, and S. Lauritzen, *Graphical Models with R*, Springer, New York, 2012.

[18] C. R. Johnson and R. L. Smith, *Path product matrices*, Linear and Multilinear Algebra, 46 (1999), pp. 177–191.

[19] C. R. Johnson and R. L. Smith, *Inverse M-matrices, II*, Linear Algebra and its Applications, 435 (2011), pp. 953–983.

[20] S. Karlin and Y. Rinott, *Classes of orderings of measures and related correlation inequalities. I. Multivariate totally positive distributions*, J. Multiv. Anal., 10 (1980), pp. 467–498.

[21] ——, *Total positivity properties of absolute value multinormal variables with applications to confidence interval estimates and related probabilistic inequalities*, Ann. Statist., 9 (1981), pp. 1035–1049.

[22] ——, *M-matrices as covariance matrices of multinormal distributions*, Linear Algebra and its Applications, 52 (1983), pp. 419 – 438.
[23] S. L. Lauritzen, *Graphical Models*, Clarendon Press, Oxford, United Kingdom, 1996.
[24] W. Ledermann, *I.—On a problem concerning matrices with variable diagonal elements.*, Proceedings of the Royal Society of Edinburgh, 60 (1940), pp. 1–17.
[25] Z. Q. Luo and P. Tseng, *On the convergence of the coordinate descent method for convex differentiable minimization*, Journal of Optimization Theory and Applications, 72 (1992), pp. 7–35.
[26] D. M. Malioutov, J. K. Johnson, and A. S. Willsky, *Walk-sums and belief propagation in Gaussian graphical models*, The Journal of Machine Learning Research, 7 (2006), pp. 2031–2064.
[27] B. F. Malle and L. M. Horowitz, *The puzzle of negative self-views: An exploration using the schema concept.*, Journal of Personality and Social Psychology, 68 (1995), p. 470.
[28] C. M. Newman, *A general central limit theorem for FKG systems*, Comm. Math. Phys., 91 (1983), pp. 75–80.
[29] A. Ostrowski, *Über die Determinanten mit überwiegender Hauptdiagonale*, Commentarii Mathematici Helvetici, 10 (1937), pp. 69–96.
[30] J. G. Propp and D. B. Wilson, *Exact sampling with coupled Markov chains and applications to statistical mechanics*, Random Structures and Algorithms, 9 (1996), pp. 223–252.
[31] C. Semple and M. A. Steel, *Phylogenetics*, vol. 24, Oxford University Press, 2003.
[32] N. Shiers, P. Zwiernik, J. Aston, and J. Q. Smith, *The correlation space of Gaussian latent tree models and model selection without fitting*, Biometrika, 103(3) (2016), pp. 531–545.
[33] M. Slawski and M. Hein, *Estimation of positive definite M-matrices and structure learning for attractive Gaussian Markov random fields*, Linear Algebra and its Applications, 473 (2015), pp. 145–179.
[34] C. Spearman, *The abilities of man*, Science, 68 (1928), p. 38.
[35] T. P. Speed and H. T. Kiiveri, *Gaussian Markov distributions over finite graphs*, Ann. Statist., 14 (1986), pp. 138–150.
[36] C. Uhler, *Geometry of maximum likelihood estimation in Gaussian graphical models*, Ann. Statist., 40 (2012), pp. 238–261.
[37] N. Wermuth and E. Scheidt, *Algorithm AS 105: Fitting a covariance selection model to a matrix*, Journal of the Royal Statistical Society. Series C (Applied Statistics), 26 (1977), pp. pp. 88–92.
[38] P. Zwiernik, *Semialgebraic Statistics and Latent Tree Models*, no. 146 in Monographs on Statistics and Applied Probability, Chapman & Hall, 2015.

Department of Mathematical Sciences, University of Copenhagen, Copenhagen, Denmark

E-mail address: lauritzen@math.ku.dk

Laboratory for Information and Decision Systems, and Institute for Data, Systems and Society, Massachusetts Institute of Technology, Cambridge, MA, USA

E-mail address: cuhler@mit.edu

Department of Economics and Business, Universitat Pompeu Fabra, Barcelona, Spain

E-mail address: piotr.zwiernik@upf.edu